Meshes preserving minimum feature size

Greg Aloupis[1], Erik D. Demaine[2], Martin L. Demaine[2], Vida Dujmović[3], and John Iacono[4] ⋆

[1] Université Libre de Bruxelles. aloupis.greg@gmail.com
[2] Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. {edemaine,mdemaine}@mit.edu
[3] Carleton University, Ottawa, Ontario, Canada. vida@cs.mcgill.ca
[4] Polytechnic Institute of New York University, Brooklyn, New York, USA. jiacono@poly.edu

Abstract. The *minimum feature size* of a planar straight-line graph is the minimum distance between a vertex and a nonincident edge. When such a graph is partitioned into a mesh, the *degradation* is the ratio of original to final minimum feature size. For an n-vertex input, we give a triangulation (meshing) algorithm that limits degradation to only a constant factor, as long as Steiner points are allowed on the sides of triangles. If such Steiner points are not allowed, our algorithm realizes $O(\lg n)$ degradation. This result[5] answers a 14-year-old open problem by Bern, Dobkin, and Eppstein.

1 Introduction

Meshing is a field frequently studied in the context of computational geometry; see [BE95,She04] for surveys. In two dimensions, the typical forms of input are point sets, polygons, and most generally, planar straight-line graphs (PSLG). The typical desired output is a decomposition into triangles or quadrangles, usually with Steiner points allowed (though their number is usually minimized as much as possible). A wide variety of quality measures dictate the desired decomposition. Often, decompositions are constructed so that there are no large angles, or instead no small angles, short edges, or short triangle heights. Most of these problems have been solved in the best sense possible. This paper highlights one problem that has not been fully solved.

⋆ Research supported by NSF grants CCF-1018370, CCF-0430849 and by an Alfred P. Sloan fellowship. Research partially completed as a visiting professor at MADALGO (Center for Massive Data Algorithmics, a Center of the Danish National Research Foundation), Department of Computer Science, Aarhus University, IT Parken, Åbogade 34, DK-8200 Århus N, Denmark.

[5] The research presented in this paper was initiated at the 24th Annual Winter Workshop on Computational Geometry at the Bellairs Research Institute of McGill University, at which the presence of Ferran Hurtado was appreciated and is hereby acknowledged.

We study the problem computing the triangulation G of a polygon P, with the possible aid of Steiner vertices. Our goal is to avoid introducing small distances between vertices and non-incident edges in G, compared to distances already existing in P. The minimum such distance in G (or for that matter in any PSLG) is called the *minimum feature size*, denoted by $\mathrm{mfs}(G)$. See [Rup93,Dey07,HMP06,Eri03]. We call the ratio $\frac{\mathrm{mfs}(P)}{\mathrm{mfs}(G)}$ the *degradation* of the decomposition of P into G. Since mfs applies to any PSLG, when it comes to measuring distances for polygons, no distinction is made between interior and exterior.

Minimum feature size is a parameter well suited for describing the resolution needed to visually distinguish elements in a mesh. For example, it measures the maximum thickness that the edges in a mesh can be drawn, while still allowing one to distinguish distinct components. Also, mfs measures the amount of error allowed in the placement of vertices, so that a drawing preserves its topology. This can be useful in manufacturing, as well as in finite element simulation.

One important issue here is the type of desired triangulation. This choice has a large effect on the results that can be achieved. See Figure 1. The most common decomposition of a polygon is the *classic triangulation*, where noncrossing chords are added between vertices of P, until the interior of P is partitioned into triangles. If we allow Steiner points, a *proper triangulation* is such that any two edges that lie on the same interior face and are incident to a common vertex are not collinear. A *nonproper triangulation* simply partitions P into triangles, with no restrictions.

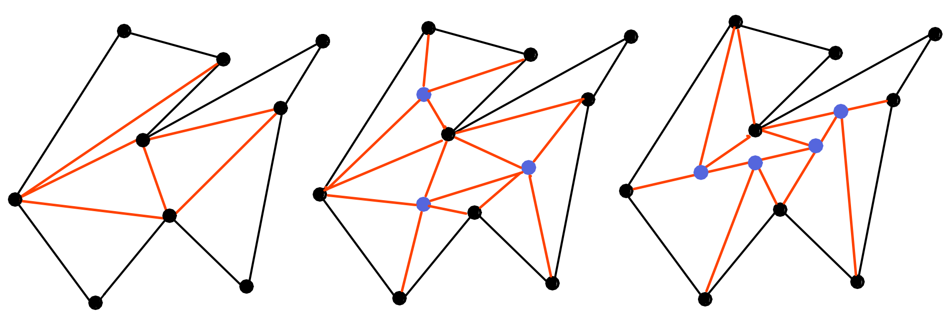

Fig. 1. Triangulation types: classic, proper, nonproper. Steiner points are blue.

Steiner points are necessary to obtain a degradation smaller than a linear factor. Specifically, Bern and Eppstein [BE95] showed that all classic triangulations of a regular n-gon have a degradation of $\Omega(n)$. This can be extended trivially to quadrangles or any decomposition with constant-size faces.

When studying this problem, Bern, Dobkin, and Eppstein [BDE95] applied the notion of *internal feature size* $\mathrm{ifs}(P)$, which is the minimum distance *inside* P between a vertex and a nonincident edge[6]. They proved that every polygon P (possibly with holes) has a nonproper triangulation in which every triangle has height $\Omega(\mathrm{ifs}(P))$. This is equivalent to guaranteeing the same bound for the ifs of the triangulation itself, and in other words means that $O(1)$ degradation is achievable for the interior of P. However, this process can reduce the minimum feature size (externally). Consequently, the first open problem the authors list is whether their result can be generalized to planar straight-line graphs.

In fact, Ruppert's Delaunay mesh refinement algorithm had already claimed constant degradation for proper triangulation of planar straight-line graphs [Rup93, Theorem 1],[7] but the "constant" actually depends on the minimum angle of the input graph (as well as the minimum triangle angle guaranteed by the algorithm).

We answer the open problem in [BDE95], by providing an algorithm that constructs a non-proper triangulation for any given polygon so that degradation is $\mathcal{O}(1)$. Our algorithm uses $\mathcal{O}(n)$ Steiner points and hence $\mathcal{O}(n)$ triangles (§2, Thm.1). Note that the triangulation of any PSLG can easily be decomposed into the triangulation of polygons. The degradation of one polygon does not affect that of another, if feature size is also measured externally. The external feature size is the ingredient that allows us to make the jump from polygons to PSLG.

Furthermore, we show that the distinction between proper and nonproper is critical. By disallowing Steiner points along the sides of triangles, our slightly modified construction gives a degradation bound of $\mathcal{O}(\lg n)$ (see §3, Thm.2). We claim that this bound is optimal, but refer the reader to a companion paper in progress (see [ADD$^+$11], which also contains a preliminary version of our upper bound).

2 Nonproper Triangulations can Preserve Minimum Feature Size

In this section we show how to construct a nonproper triangulation for any polygon P, such that the degradation is $\Theta(1)$. We use $\Theta(n)$ Steiner points, and the construction can be computed in linear time.

We begin by explaining how to triangulate parallelograms and trapezoids. Trivially any rectangle can be triangulated by placing a Steiner vertex s at its center. The mfs will degrade by a factor of 2. Suppose instead that we are given

[6] Note that "internal feature size" is called "minimum feature size" in [BDE95].
[7] Incidentally, this is also the paper that first introduced the notion of feature size.

a parallelogram P with top and bottom edges horizontal, and tilted towards the right. A segment with one endpoint at the lower-right vertex determines $\mathrm{mfs}(P)$. Its direction depends on the height of P and the length of the horizontal edges, The segment is either vertical representing the height, or orthogonal to the left edge of P. Either way, placing s at the center yields a degradation of 2, as can be easily verified by examining similar parallelograms. Specifically the new mfs will be determined by a segment parallel to the original one, from s to the boundary of P.

Lemma 1. *Any trapezoid T can be triangulated with a degradation $d_{trap} \leq 2$.*

Proof. Let L be the shorter of the parallel edges on T, with length ℓ, and *wlog* at the bottom of T. Let U be the top edge of T, and h be the height of T. Consider the rectangle R obtained by projecting L vertically upward onto the line through U. Suppose that R is contained in T. Then $\mathrm{mfs}(T) = min\{h, \ell\}$ (i.e. it is determined by the dimensions of R). We place a Steiner vertex s at the middle of R. The distance from s to L or U is $\frac{h}{2}$. The distance from s to the sides of T is greater than $\ell/2$. So if $h \leq \ell$, the degradation of the resulting triangulation is 2. Otherwise it is even less.

Now suppose that R is not contained in T, in which case we know that both side edges of T are slanted in the same direction, *wlog* towards the right. Then $\mathrm{mfs}(T)$ is determined by a segment with one endpoint on the lower-right vertex of T. Consider the parallelogram P obtained by sweeping a horizontal segment of length ℓ from L to U, while keeping its left endpoint on the left side of T. As before, T and P have the same minimum feature size determined by the same segment. The center of P is suitable for s. By construction, s separates P into four similar parallelograms. By preceding arguments described for parallelograms, the degradation of the resulting triangulation is 2. □

Lemma 2 (Perturbation Lemma). *Moving all the vertices of a PSLG G by at most $\alpha\,\mathrm{mfs}(G)$, for $\alpha < \frac{1}{2}$, results in a PSLG G' with degradation at most $\frac{1}{1-2\alpha}$ relative to G. The drawings of G and G' are combinatorially equivalent.*

Proof. Any distance determined by a point and a non-incident edge can be shortened by at most $\alpha\,\mathrm{mfs}(G)$ at each end, and thus $2\alpha\,\mathrm{mfs}(G)$ total. These distances were at least $\mathrm{mfs}(G)$ to begin with. Combinatorial equivalence follows from the fact that no vertex is allowed to move enough to cross a non-incident edge. □

The next lemma is essentially the most critical element of the main theorem that will follow.

Lemma 3. *Let R be a rectangle with $\mathrm{mfs}(R) = s$, width equal to a multiple of s, and height $h > s$. Allow Steiner vertices to be placed along the top edge of R at any selected locations that are multiples of s from one of its endpoints. Then R can be non-properly triangulated, without adding any additional Steiner vertices on its boundary, with constant degradation.*

Proof. The proof is by construction, specifically the triangulation G shown in Figure 2 (C), where as a worst case scenario we have placed Steiner vertices at every multiple of s on the bottom edge of R.

The main construction has a set of edges anchored at one corner of R (upper-left in the figure). Starting from the shortest (and most clockwise), each such edge e_i reaches to a horizontal coordinate twice as large as the previous one, and to a vertical coordinate $\frac{h}{3}$ from the bottom, where it meets the midpoint of a vertical edge g_i of length $\frac{2h}{3}$. Also, between every two such successive edges g_i, there is a region below one of the anchored edges, that has a sawtooth pattern matching the Steiner vertices. Specifically, for e_i the sawtooth region is bounded by e_i, g_i, g_{i-1}, and the bottom ege of R.

Among the newly constructed Steiner vertices on e_i (i.e., on the top side of its corresponding sawtooth), the leftmost, S_ℓ, is closest to edge e_{i+1}. This vertex happens to be the intersection of e_i with g_{i-1}. Since e_{i+1} reaches twice as far as e_i, but to the same vertical coordinate, the vertical separation between e_{i+1} and S_ℓ is $\frac{h}{6}$. The horizontal separation between e_{i+1} and S_ℓ is at least s. We conclude that the distance between e_{i+1} and S_ℓ is at least $s\frac{h}{6}\frac{1}{\sqrt{(\frac{h}{6})^2+s^2}}$. This will dominate the feature size of all triangles emanating from the top-left of R.

Each sawtooth consists of the bottom horizontal edge, two vertical edges (the left twice the length of the right), and a tilted top, where the tilt angle gets closer to horizontal as the sawtooths move further to the right. The minimum distance created within any sawtooth occurs at its right hand side and is determined by the angle of the internal diagonals and by the tilt of the top. The distance is minimized as the tilt of the top increases and as the diagonal becomes less vertical. So, the minimum distance overall is to be found at the leftmost sawtooth, in its rightmost component (triangle). The new distance introduced is at least $\frac{h}{3}\frac{s}{2}\frac{1}{\sqrt{(\frac{h}{3})^2+(\frac{s}{2})^2}}$. Notice that both terms calculated imply constant degradation. $\square$

We now state the main theorem of this section:

Theorem 1. *Every polygon has a non-proper triangulation with constant degradation.*

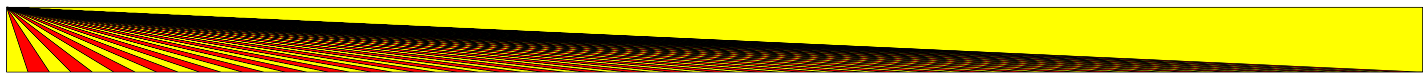

(A) A simple proper triangulation of a rectangle with 66 vertices along the bottom edge. Observe that the interiors of many of the triangles are not discernible. Note that in this figure, as well as in (B) and (C), no additional Steiner vertices have been placed on the exterior of the rectangle.

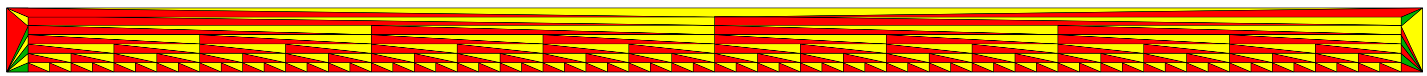

(B) Proper triangulation of the same rectangle as in (A) using a skip-list-like divide-and-conquer approach. All triangles are discernible, but as this figure is generalized to larger rectangles the minimum feature size will slowly degrade.

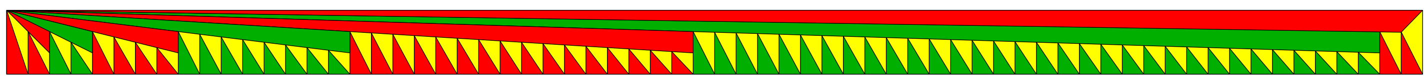

(C) Proper triangulation of the same rectangle in (A) and (B) using a novel construction. The triangulation is much easier to see than the previous two, and generalizing it to longer rectangles will not change the minimum feature size beyond that of this figure.

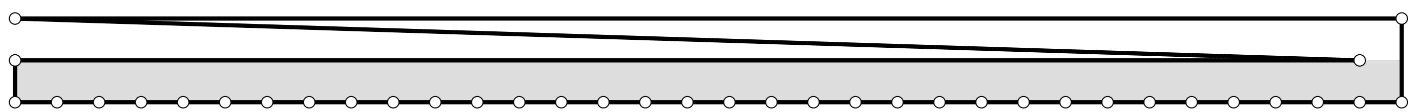

(D) This polygon illustrates that Steiner vertices cannot be placed on a significant fraction of the boundary close to the reflex vertex. A construction like one of the above three can be used to triangulate the gray shaded area.

Fig. 2.

Proof. Let P be an n-gon with minimum feature size 1. Let the curve P_2 be the locus of points inside P that have minimum distance $\frac{1}{4}$ from ∂P. This is obtained from the well-known grassfire transformation. P_2 is a closed curve consisting of n line segments (one per segment of P) as well as one circular arc corresponding to every reflex vertex of P. Each such arc spans less than $180°$. See Figure 3 for an illustration of this process. Because P has minimum feature size 1, each line segment in P_2 has length at least $\frac{1}{2}$ (see Appendix A).

P_2 splits the interior of P into two regions, which we call the *Interior* and the *Tube*. P_2 itself belongs to both regions. We will modify and refine this boundary a few times, and then triangulate each region separately. Any operations in the interior of one region will not affect degradation in the other.

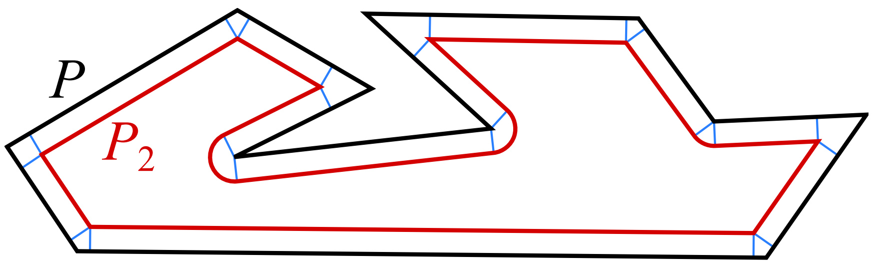

Fig. 3. A polygon P and the closed curve P_2 created by the grassfire tranformation.

Refinement of P_2: We create a polygon P_3 by replacing all circular arcs on P_2 with polylines. Let a be a given circular arc with endpoints a_1 and a_2. If a spans more than $90°$, we can replace it with segments $a_1 m$ and $m a_2$, where m is the midpoint of the arc. The minimum distance between P_3 and P is $\frac{1}{4\sqrt{2}}$ (achieved when a approaches $180°$). So far, this dominates the feature size of the Tube, since (when $a = 90°$) the segments on P_3 can have length as short as the side of a regular octagon of diameter $\frac{1}{4}$. That is, $\frac{1}{4} 2 \sin \frac{\pi}{8}$; (appx. 0.19). We say that the Tube degradation is at most $d_{\mathit{Tube3}} = 4\sqrt{2}$. On the other hand, the feature size $\frac{1}{d_3}$ of the Interior is $\frac{1}{4} 2 \sin \frac{\pi}{8}$, i.e. dominated by its boundary, P_3. Distances through the interior of P_3 are still at least $\frac{1}{2}$.

If instead a spans less than $90°$, we extend its adjacent edges on P_2, through a_1 and a_2 respectively, until they meet. This extension remains at a distance greater than $\frac{1}{4}$ from P, so the Tube degradation is unaffected. The extension also remains at most $\frac{\sqrt{2}}{4}$ from the vertex on P that generated the arc (the max is achieved when the angle is $90°$). All other points on P_3 are even closer to the boundary of P. Thus no two points from different edges on P_3 will get closer than $1 - \frac{\sqrt{2}}{2}$ to each other (roughly 0.29; not enough to dominate the feature size of the Interior).

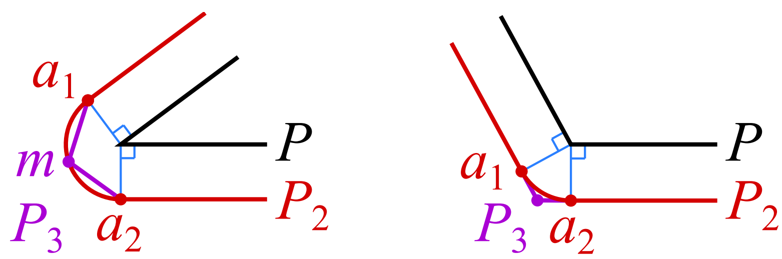

Fig. 4. How to transform P_2 to P_3. Left: arc spans more than $90°$; Right: arc spans less than $90°$.

Let P_4 be formed by snapping the vertices of P_3 vertically to a horizontal grid of granularity $g = \frac{1}{2d_3}$. Any point can snap at most a distance of $\frac{1}{4d_3}$; half the grid size. A pair of points on P_3 at a *nearly* co-vertical position a distance of $\frac{1}{d_3}$ from each other may snap towards each other, so snapping can degrade the feature size of P_3, and thus also the Interior, by a factor of 2. The effect on the minimum distance between P and P_3 is smaller, since vertices of P remain fixed. This distance can drop to $\frac{1}{d_{Tube3}} - \frac{1}{4d_3}$. At this point though, the minimum distance in both regions is to be found on their common boundary, and the value is $\frac{1}{2d_3}$.

Triangulation of Interior:

Let P_5 consist of P_4 and the horizontal trapezoidation of the interior of P_4. All vertices on P_4 lie on the grid and thus the feature size is preserved. Let P_6 consist of the triangulation of P_5 obtained by placing a vertex in each trapezoid according to the method presented in Lemma 1. Thus, the degradation in this step is $d_6 = d_{trap}$.

The total degradation of the Interior is therefore $2d_3d_6$. The feature size of the Interior is $\frac{1}{8}\sin\frac{\pi}{8}$ (appx. 0.047).

Triangulation of Tube:

Obtaining a triangulation of the Tube is done without adding more Steiner vertices to its boundary, and thus any degradation in this process will not amplify the degradation of Interior, or affect feature size via the exterior of P.

Before proceeding to the algorithm, which has four main steps, we require one definition: A *quasi-trapz* is a quadrilateral that can be transformed into a trapezoid by perturbing its vertices by an amount small enough so that the Perturbation Lemma can be applied.

1. **Subdivision of the Tube into triangles and quasi-trapz.** Consider all convex polygonal chains that were created in P_3 as replacements of circular arcs spanning more than $90°$ on P_2. Recall that such chains consist of two segments, which by now can also contain Steiner vertices from the trapezoidation of P_4. For each chain, we connect the endpoints to the unique reflex vertex v of P that generated the corresponding (replaced) arc via the grassfire transform. Similarly, we connect every convex vertex of P to its corresponding convex vertex on P_6, and we connect reflex vertices of P (with arcs spanning less than $90°$) to their unique corresponding reflex vertex on P_6. This subdivides the Tube into quasi-trapz and *convex fans* (i.e. a vertex visible from a convex chain). In fact any such fan is just a quadrilateral, since the chain opposite v had only two edges on P_3 (with Steiner vertices added later on). See Figure 5.

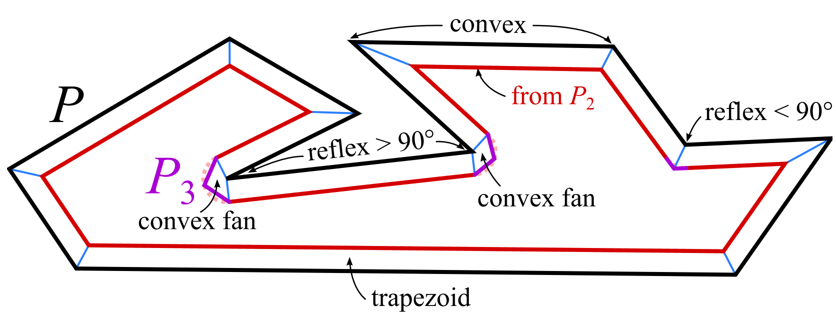

Fig. 5. Phase 1 of Tube triangulation: Subdividing Tube into fans and quasi-trapz. Angles indicated correspond to arc spans.

The only degradation caused can be due to a newly created edge, e, and some non-incident vertex p on P_6. What matters is the angle that e makes with P_6, and the proximity of p to the endpoint of e on P_6. The latter is at least $\frac{1}{2d_3}$. Without taking snapping into account, the aforementioned angle would be no less than $45°$. See Figure 6. So, adding these edges would only degrade feature size from $\frac{1}{2d_3}$ by a factor of $\sqrt{2}$. This means that the feature size of the fans could drop to roughly 0.067. This is not close to the smaller feature size in other areas that will be created, and snapping the vertices of the fan will not have any significant effect. The snapping would have to be so extreme that the angle mentioned would drop from $45°$ to under $14°$. This is not possible when two vertices of a triangle move by less than a quarter of the shortest length.

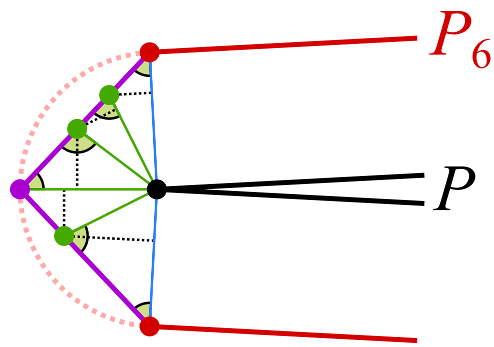

Fig. 6. Triangulating a convex fan: the angle between P_6 (purple) and edge e (blue) is at least $45°$. The same holds for other edges (green) from P to Steiner points on the fan. Dotted black segments are minimum distances in newly created triangles.

We continue to triangulate each fan by adding diagonals from v to all remaining vertices within. The preceding analysis follows verbatim. Structurally, the end result is that any non-triangulated region of the Tube is a quasi-trapz that has an edge of P (the *bottom*) and an edge of P_6 (the *top*) on its boundary. Were it not for the snapping, the top and bottom would be parallel, and the quasi-trapz would be a trapezoid. Note that the top can contain Steiner vertices, generated during the trapezoidation of P_4.

2. **Subdivision of quasi-trapz with Steiner vertices on the top.** We will subdivide quasi-trapz so that the only remaining non-triangulated regions will be quasi-trapz without Steiner points on their boundary, and rectangles possibly with such Steiner points.
Let Q be a quasi-trapz to be subdivided. By assumption, Q has at least one Steiner vertex on its top, T. We will start adding diagonals from the bottom, B, to Steiner vertices on T. This will progressively cut off triangles, leaving a smaller quasi-trapz Q'. As we do this, we will be shortening T, so that it either has no Steiner vertices on it, or the internal angles of Q' at T are at least $135°$. This will also imply that T comfortably projects onto the bottom, B, in a direction orthogonal to T.
For each endpoint t of T with internal angle smaller than $135°$ (note that the angle is at least $45°$ to start with), do the following. Let v be the neighbor of t on Q, on edge B. Traverse T from t until a Steiner vertex s is found, and join s to v. While the angle condition is not met, keep forming such an

edge to v for each successive s. This repeatedly cuts off triangles from Q, until it is either a triangle or Steiner-free quasi-trapz (if no more Steiner vertices remain), or until the angle condition is satisfied. The triangles cut off on each side form a convex fan triangulation (with v as the apex), identical in nature to those described in phase 1. Thus the degradation caused so far can be absorbed into the preceding analysis. See Figure 7.

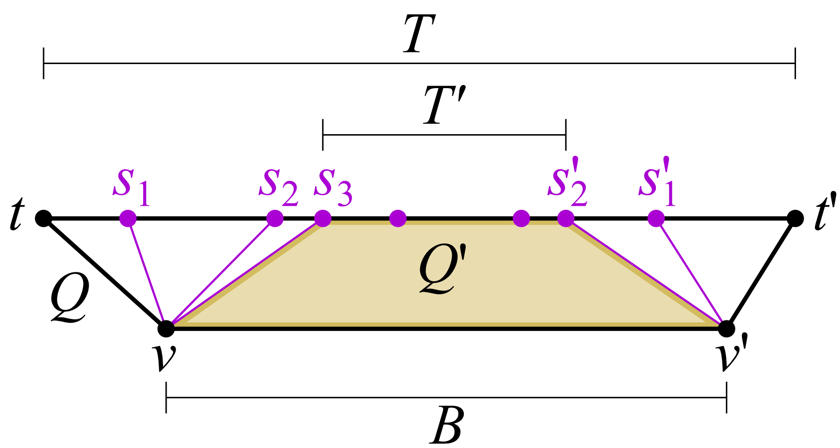

Fig. 7. Phase 2 of Tube triangulation: decomposing Q into Q' and triangulated fans. Either the internal angles of Q' at the top are greater than $135°$, or Q' has no Steiner vertices.

Now we can assume that we have a subset Q' of Q, with the desired projection of T onto B. We consider the shape of Q' as it existed before snapping, along with corresponding subsets of T and B, labeled T' and B' respectively. Recall that before snapping, T and B were parallel, at distance $\frac{1}{4}$. The minimum feature size of Q' is dominated by the separation of Steiner vertices on its boundary, i.e. $\frac{1}{2d_3}$. Let a and b be the endpoints of T. Let a' and b' be vertices placed at a distance $\frac{1}{4} - \frac{1}{2d_3}$ away from a and b, respectively, so that $baa'b'$ is a rectangle R inside Q'. Notice that a' and b' are $\frac{1}{2d_3}$ from B, and even further from the sides of Q'. So this placement doesn't affect the feature size. Now, connect a' and b' to the vertices on Q'. Inside Q', we are left with a Steiner-free trapezoid Q'' (it is below $a'b'$ and will again become a quasi-trapz when we account for snapping), the rectangle R with Steiner points on the edge ab, and two triangles. See Figure 8.

Finally we must reinstate the snapping of the segment ab. The positions of a' and b' will follow so that R moves rigidly. We now examine the effect

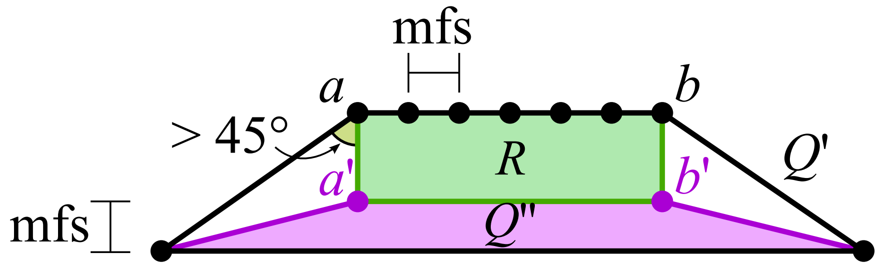

Fig. 8. Phase 2 of Tube triangulation: decomposing Q' into R and Q''.

of the motion of a' on the feature size of the components of Q'. Recall that a' is snapped by at most $\frac{1}{4d_3} = \frac{1}{16\sqrt{2}}$ (roughly 0.044). So it can reduce the feature size of Q'' to $\frac{1}{4d_3}$ (by moving half way to the fixed edge B on Q''. The effect of a' is even smaller on the feature size of the triangles in Q', since its distance to their non-incident edges is greater than its distance to B. Of course, there is no effect on R.

3. **Triangulate all remaining Steiner free quasi-trapz.** Here we triangulate any quasi-trapz Q'' constructed in the previous phase. One Steiner vertex s suffices, as with any trapezoid. There is a placement for s in the corresponding unsnapped trapezoid Q so that degradation is no more than 2. Because Q has height $(\frac{1}{2d_3})$, s would normally be placed between the parallel edges of Q, i.e. $\frac{1}{4d_3}$ from B. However the edge $a'b'$ might snap by this much, and this would create an arbitrarily small distance to s. So, instead we will place s at a distance $\frac{1}{8d_3}$ from B, since B will remain fixed. The effect is that the feature size of Q'' can be reduced to $\frac{1}{8d_3}$, but no less. Currently, this dominates the overall feature size, at roughly 0.022. See Figure 9.

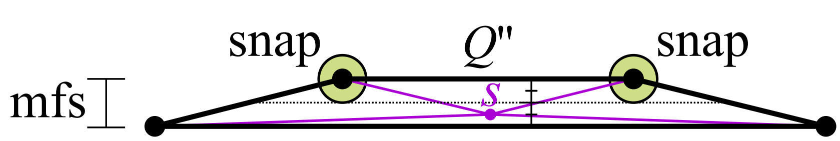

Fig. 9. Phase 3 of Tube triangulation: handling the last remaining quasi-trapz.

4. **Triangulate rectangles generated in phase 2.** For each rectangle $R = abb'a'$, we use the construction presented in Lemma 3. R has height $h = \frac{1}{4} - \frac{1}{2d_3}$, and Steiner vertices are placed on one of its longer sides, at dis-

tances of $\frac{1}{2d_3}$. Then the formula in Lemma 3 yields a value of greater than 0.024 for $\mathrm{mfs}(R)$.

The important conclusion is that each step described incurs a constant degradation, therefore the aggregate is also constant. Most of the steps described probably have tighter bounds. Furthermore, these steps could be optimized to work more harmonically, or replaced with more efficient constructions. For the record, we claim here that the degradation is under 45 (the inverse of 0.022 calculated in phase 3 in the Tube). $\square$

3 Degradation Upper Bound for Proper Triangulations

Theorem 2. *Every n-vertex polygon has a proper triangulation with $\mathcal{O}(\log n)$ degradation.*

Proof. The algorithm of the previous section can be used, with the exception of the non-proper triangulation of the rectangle $R = abb'a'$ that must be replaced. That triangulation is now done using the construction of Figure 2 (B). This simple skip-list like decomposition of a rectangle into proper triangles has degradation of $\mathcal{O}(\log n)$.

4 Discussion

Although several steps in our construction require subtle constructions and details to keep things at constant distance, we believe that the essential hurdle was how to triangulate rectangles with many Steiner points on a side, without adding new Steiner points on the boundary. This was the breakthrough needed to solve the problem at hand.

Preserving minimum feature size is not the only priority in meshing, but it is a meaningful measure of mesh quality. We leave to future work the possibility of simultaneously attaining small degradation and other mesh properties such as maximum angle bounded away from $180°$. This may be possible by simply combining algorithms.

An application of our result exists in the 197-year-old algorithm by Lowry [Low14,Fre97] for finding a common dissection of two given polygons of equal area. This algorithm starts by triangulating the polygon P, then uses a dissection from 1778 to convert each triangle into a rectangle with a common height ε equal to half the minimum height of all triangles. The algorithm uses $\mathcal{O}(A/\varepsilon)$ pieces where A is the area of the polygon, and this bound is clearly the best possible for each individual piece. The issue is that ε depends on the choice of triangulation; without care (as in all previous descriptions of this algorithm), a triangulation could have an arbitrarily small triangle height, and thus force an arbitrarily small ε. If we use the triangulation algorithm by Bern *et al.* [BDE95] we obtain a dissection using $\mathcal{O}(A/\operatorname{ifs}(P)) = \mathcal{O}(r^2)$ pieces. Our results generalize to polygons that have already been partially dissected, where we obtain a bound of $\mathcal{O}((A/\operatorname{mfs}(P))) = \mathcal{O}(r^2 n)$ pieces. These are the first pseudopolynomial bounds on dissection.

Finally, we see no reason why our method cannot be used to create a proper triangulation with constant degradation of *internal* feature size. Recall that this has already been done for non-proper triangulation [BDE95]. To obtain constant degradation internally and externally, we were forced to introduce some non-properly triangulated component (recall our claim that with only proper triangulation, constant degradation is impossible). However, the only component where our triangulation is non-proper is within the quasi-trapz in the Tube (and specifically this unavoidable component was pushed further into rectangular regions). Instead, if we are not concerned with external feature size, we can triangulate the quasi-trapz by creating Steiner vertices on the boundary of P, to match those on the inner boundary of the Tube. This would work just as well for polygons with holes, as does our entire method.

References

[ADD$^+$11] Greg Aloupis, Erik Demaine, Martin Demaine, Vida Dujmovic, and John Iacono. Meshes preserving minimum feature size. Technical report arXiv:0908:2493, 2011.

[BDE95] Marshall Bern, David Dobkin, and David Eppstein. Triangulating polygons without large angles. *Interational Journal of Computational Geometry & Applications*, 5(1–2):171–192, March–June 1995.

[BE95] Marshall Bern and David Eppstein. Mesh generation and optimal triangulation. In Ding-Zhu Du and Frank Kwang-Ming Hwang, editors, *Computing in Euclidean Geometry*, number 4 in Lecture Notes Series on Computing, pages 47–123. World Scientific, second edition, 1995.

[Dey07] Tamal K. Dey. Delaunay mesh generation of three dimensional domains. Technical Report OSU-CISRC-09/07-TR64, Ohio State University, October 2007.

[Eri03] Jeff Erickson. Nice point sets can have nasty delaunay triangulations. *Discrete & Computational Geometry*, 30(1), July 2003.

[Fre97] Greg N. Frederickson. *Dissections: Plane and Fancy*. Cambridge University Press, November 1997.

[HMP06] Benoit Hudson, Gary Miller, and Todd Phillips. Sparse Voronoi Refinement. In *Proceedings of the 15th International Meshing Roundtable*, pages 339–356, Birmingham, Alabama, September 2006. Springer-Verlag.

[Low14] Mr. Lowry. Solution to question 269, [proposed] by Mr. W. Wallace. In T. Leybourn, editor, *Mathematical Repository*, volume 3, part 1, pages 44–46. W. Glendinning, London, 1814.

[Rup93] Jim Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Proceedings of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 83–92, Austin, Texas, 1993.

[She04] Jonathan Richard Shewchuk. Theoretically guaranteed Delaunay mesh generation—in practice. Short course at 13th International Meshing Roundtable, 2004. `http://www.cs.berkeley.edu/~jrs/papers/imrtalk.pdf`.

Appendix A

Claim: *Let C be the curve that is the locus of points inside polygon P, at distance $\frac{1}{4}$ from the boundary of P. If $\mathrm{mfs}(P) = 1$ then every straight edge of C has length at least $\frac{1}{2}$.*

C is obtained from P using the grassfire transform, commonly used as a visualization of the construction of the medial axis. Note that each edge e on P transforms to an edge e' on C continuously as the grassfire progresses. The shape of e' depends on local conditions; specifically the angles of vertices at the endpoints of e. The edge e' must reside on the horizontal line at a distance $\frac{1}{4}$ below e.

Let e be positioned horizontally, between vertices p_1, p_2. If both endpoints are reflex vertices on P, then e' will have the same length. If one of the vertices is reflex (*wlog* the left, p_1), then the left endpoint v_1 of e' will be located vertically below p_1, at a distance of $\frac{1}{4}$. Follow a ray to the right of v_1 for a distance of $\frac{1}{2}$, to construct a point, x. Let Q be the unit quarter-circle in the lower-right quadrant of p_1. Then x is inside Q and at a distance greater than $\frac{1}{4}$ from the arc of Q. Thus x cannot be a vertex on C, since it is not at a distance $\frac{1}{4}$ from any point on P (excluding e itself). This implies that e' has length greater than $\frac{1}{2}$.

Finally, there is the case where both endpoints of e are convex vertices. Note that e' can have a length of $\frac{1}{2}$ if e has length 1 and both convex angles are $90°$. Then, the endpoints of e' are directed inward at angles of $45°$, relative to the endpoints of e. If $|e| = 1$, then both convex angles must be at least $90°$, so $|e'| \geq \frac{1}{2}$.

We can make e larger to allow for smaller convex angles at its endpoints. If we do so, the worst scenario is one where the edges adjacent to e in P are angled in a way that they eventually have a distance of 1 with each other (i.e., we close the angles as much as possible without violating feature size of P). So, *wlog* assume that the angle at p_1 is less than $90°$. Follow the edge u neighboring e to the left until hitting a horizontal line at a distance of 1 from e. Note that this intersection point, y, must exist, otherwise we contradict the mfs assumption about P (a vertex of e' would be too close to e). In other words y belongs to u. Construct the point z at a distance 1 vertically above y. This must be part of e otherwise again we contradict the mfs assumption about P (there would be a vertex to the left of z, too close to u). No part of the boundary of P intersects the triangle yze_1. The left endpoint t_1 of e' cannot be more than $\frac{1}{4}$ to the right of the segment yz; this happens if the angle at p_1 is $90°$, and as it is opened t_1 moves to the left relative to yz.

Consider the unit quarter-circle Q centered at y, in its top-right quadrant. No part of P can intersect Q; if an edge not adjacent to y does so, it will be too close to y, and if an adjacent edge does so (if y were a real vertex), it will be too close to e. Now consider the vertex p_2, common to e and its edge v to the right. Wherever p_2 is, $|e'|$ will be minimized if we minimize the angle at p_2. If this angle is less than $90°$, then since v must miss Q, we follow the same reasoning as above to easily conclude that the e' has length greater than 1. That is, the right end t_2 of e' will not be more than $\frac{1}{4}$ to the left of some vertical segment $y'z'$ analogous to yz.

So we are left with the case where p_2 has angle greater than $90°$ and is located within one unit of z, i.e. its vertical projection intersects Q. As mentioned, $|e'|$ will be minimized if the angle at p_2 is minimized, which is to say that v should be rotated as clockwise as possible, until it becomes tangent to Q. This means that the bisector at p_2 intersects y. So p_2 should be placed as close to z as possible, to minimize $|e'|$. We will now work within the triangle p_1yp_2.

Let z' be the intersection of yz with the line through e'. We know that t_2 will be to the right of z'. On the other hand, t_1 can be arbitrarily to the left, or slightly to the right (specifically no more than $\frac{1}{4}$). Let $t_2 - z$ be e'_2, and $z - t_1$ be e'_1. Then $|e'|$ is equal to $e'_1 + e'_2$.

There is another constraint on the location of p_2: it cannot be too close to u. Let α be the angle at p_1. Then $\sin\alpha = \frac{1}{|e|}$, meaning p_2 is placed so that its distance to u is 1. Let $z - p_1$ be e_1. From triangle yzp_1 we have $\tan\alpha = \frac{1}{e_1}$.

Let e_2 be $p_2 - z$. Triangle yzp_2 is similar to triangle $yz't_2$, so $e'_2 = \frac{3e_2}{4}$. From the triangle formed by p_1, t_1, and the projection of t_1 onto e, we have $\tan\frac{\alpha}{2} = \frac{0.25}{e_1 - e'_1}$; ($t_1$ is $\frac{1}{4}$ below e, and on the bisector of α).

Reordering and combining from above, we have:
$e'_1 = e_1 - \frac{1}{4\tan\frac{\alpha}{2}} = \frac{1}{\tan\alpha} - \frac{1}{4\tan\frac{\alpha}{2}}$. Since $e_2 = \frac{1}{\sin\alpha} - e_1$, we have
$e'_2 = \frac{3}{4}(\frac{1}{\sin\alpha} - e_1) = \frac{3}{4}(\frac{1}{\sin\alpha} - \frac{1}{\tan\alpha})$.

Notice that when α increases to $90°$ (i.e. p_1 approaches z from the left), we have $e'_1 = -\frac{1}{4}$ and $e'_2 = \frac{3}{4}$, so $|e'| = \frac{1}{2}$. For $\alpha \leq 90°$, e' is longer.